charles river analytics

Charles River Analytics Inc.
625 Mount Auburn St.
Cambridge, MA 02138
www.cra.com

p: 617.491.3474
f: 617.868.0780

20 November 2013

Dr. Harold Hawkins
ONR Code 341
Office of Naval Research
875 North Randolph SL
Arlington. VA 22203-1995

Reference: US Navy Contract N00014-12-C-0653: “The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation”
Charles River Analytics Contract No. C12186

Subject: Contractor’s Quarterly Status Report #5
Reporting Period: 20-August-2013 to 19-November-2013

Dear Dr. Hawkins,

Please find enclosed 1 copy of the Quarterly Status Report for the referenced contract. Please feel free to contact me with any questions regarding this report or the status of the “The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation” effort.

Sincerely,

W. Scott Neal Reilly

W. Scott Neal Reilly
Principal Investigator

cc: Michael Hession, DCMA
Annetta Burger, ONR
Whitney McCoy, Charles River Analytics

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **10 NOV 2013**	2. REPORT TYPE	3. DATES COVERED **00-00-2013 to 00-00-2013**

4. TITLE AND SUBTITLE **The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Charles River Analytics,625 Mount Auburn Street,Cambridge,MA,02138**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**20**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

Charles River Analytics

Monthly Technical Progress Report No. R12186-05
Reporting Period: August 20, 2013 to November 19, 2013
Government Contract No. N00014-12-C-0653
Charles River Analytics Contract No. C12186

The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation

Quarterly Status Report

Principal Investigator: Scott Neal Reilly

Charles River Analytics
625 Mount Auburn Street
Cambridge, MA 02138
617-491-3474

November 20, 2013

1. Executive Summary

The proposed research effort builds on and extends the work of the previous ONR-funded "Validation Coverage Toolkit for HSCB Models" project. The overall objectives of the on-going research program are:

- Help scientists create, analyze, refine, and validate rich scientific models
- Help computational scientists verify the correctness of their implementations of those models
- Help users of scientific models, including decision makers within the US Navy, to use those models correctly and with confidence
- Use a combination of human-driven data visualization and analysis, automated data analysis, and machine learning to leverage human expertise in model building with automated analyses of complex models against large datasets

Specific objectives for the current effort include:

- **Fluid temporal correlation analysis.** Our objective is to design a new method for performing temporally fluid correlation analysis for temporal sets of data and implement the method as a new prototype component within the Model Analyst's Toolkit (MAT) software application.
- **Automated suggestions for model construction and refinement.** Our objective is to design and implement a prototype mechanism that learns from data how factors interact in non-trivial ways in scientific models.
- **Data validation and repair.** Our objective is to design and implement a prototype capability to identify likely errors in data based on anomalies relative to historic data and to use models of historic data to offer suggested repairs.
- **System prototyping.** Our objective is to incorporate all improvements into the MAT software application and make the resulting application available to the government and academic research community for use in scientific modeling projects.
- **Evaluation of applicability to multiple scientific domains.** Our objective is to ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains by identifying and building at least one neurological and/or physiological model and analyze the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model.

2. Overview of Problem and Technical Approach

2.1. Summary of the Problem

One of the most powerful things scientists can do is to create models that describe the world around us. Models help scientists organize their theories and suggest additional experiments to run. Validated models also help others in more practical applications. For instance, in the hands of military decision makers, human social cultural behavior (HSCB) models can help predict instability and the socio-political effects of missions, whereas models of the human brain and

mind can help educators and trainers create curricula that more effectively improve the knowledge, skills, and abilities of their pupils.

While there are various software tools that are used by the scientific community to help them develop and analyze their models (e.g., Excel, R, Simulink, Matlab), they are largely so general in purpose (e.g., Excel, R) or so focused on computational models in particular (e.g., Simulink, Matlab), that they are not ideal for rapid model exploration or for use by non-computational scientists. They also largely ignore the problem of validating the models, especially when the models are positing causal claims as most interesting scientific models do. To address this gap, Charles River Analytics undertook the "Validation Coverage Toolkit for HSCB Models" project with ONR. Under this effort, we successfully designed, implemented, informally evaluated, and deployed a tool called the Model Analyst's Toolkit (MAT), which focused on supporting social scientists to visualize and explore data, develop causal models, and validate those models against available data (Neal Reilly, 2010; Neal Reilly, Pfeffer, Barnett et al., 2011, 2010).

As part of the development of the MAT tool, we identified four important extensions to that research program that would further support the scientific modeling process:

- Correlation analyses are still the standard way of identifying relationships between factors in a model, but correlations are fundamentally flawed as a tool for analyzing potentially causal or predictive relationships as they assume instantaneous effects. Even performing correlation analyses with a temporal offsets between streams of data is insufficient as the temporal gap between the causal or predictive event and the following event may not be the same every time (either because of variability in the system being modeled or because of variability introduced by a fixed sampling rate). What we need is a novel way of evaluating the true predictive power across streams of data that can deal with fluid offsets between changes in one stream of data and follow events in the other stream of data.
- Modeling complex phenomena is a fundamentally difficult task. Human intuition and analysis is by far the most effective way of performing this task, but even humans can be overwhelmed by the complexity of modeling the systems they are studying (e.g., socio-political system, human neurophysiology). Automated tools, while not especially good at generating reasonable scientific hypotheses, *are* extremely good at processing large amounts of data. We believe there is an opportunity for computational systems to enhance human scientific inquiry. Under the "Validation Coverage Toolkit for HSCB Models" project, we demonstrated how automated tools could help human scientists to analyze and validate their models against data. We believe a similar approach can be used to help suggest modifications to the human-built models to make them better match the available data. To be useful, however, such automated analyses will need to be rich enough to suggest subtle data interactions that are most likely to be missed by the human scientist. For instance, correlations (especially correlations that take into account fluid temporal displacements) could be used to identify likely relationships between streams of data, but such an approach would miss complex, non-linear relationships between interrelated factors that cannot be effectively analyzed with

simple two-way correlations. For instance, if crime waves are associated with increases in unemployment *or* drops in the police presence, that would be hard to identify with a correlation analysis. We need richer automated data analysis techniques that can extract complex, non-linear, multi-variable relationships between data if we are to effectively suggest model improvements to human scientists.

- Even if a scientific model is sound, if the data sets provided as inputs to the model are unreliable, the results of the model are still suspect. And, unfortunately, data will often be wrong. For instance, HSCB surveys are notoriously unreliable and biased for a variety of reasons, and neurological and physiological data can be corrupted by broken or improperly used sensors. If it were possible to identify when data was unreliable and, ideally, even repair the data, then the models that are using the data could once again be effectively used.
- The MAT tool we developed under the "Validation Coverage Toolkit for HSCB Models" project was focused primarily on assisting social scientists in the analysis, refinement, and validation of HSCB models. In parallel with that effort, however, we also took an opportunity to apply MAT to evaluating neurological and physiological data under the DARPA-funded CRANIUM (Cognitive Readiness Agents for Neural Imaging and Understanding Models) program. We discovered the generality of the MAT tool makes it potentially applicable to a great number of different scientific domains. MAT proved to be a useful, but peripheral tool, in CRANIUM. We believe MAT could be applied to a broader suite of scientific modeling problems than it has been so far.

2.2. Summary of our Approach

To address these identified gaps and opportunities, we are extending MAT's support for model development, analysis, refinement, and validation; enhancing MAT to analyze and repair data; and demonstrating MATs usefulness in additional scientific modeling domains. Our approach encompasses the following four areas, which correspond to the four gaps/opportunities identified in the previous section:

- **Temporally Fluid Correlation Analysis.** We are designing a new method to perform Temporally Fluid Correlational Analysis on temporal sets of data, and we are implementing the method as a new component within the MAT software application. The version of MAT at the beginning of the new effort supported correlation analysis for temporally offset data; it shifts the two data streams being compared by a fixed offset that is based on the sampling rate of the data (i.e., data that is sampled annually will be shifted by one year at a time), performs a standard correlation on the shifted data, plots the correlation value against the amount of the offset, and then repeats the process for the next offset amount. If two data streams are shifted by a fixed offset (e.g., changes in one stream are always followed by a comparable value in the other stream after a fixed time), then this method will find that offset. Under the current effort, we are expanding on this capability to support fluid temporal shifts within the data streams. That is, we are making it possible to identify when the temporal offset between the

change in the first data stream and its effect in the second stream is not a static amount of time.

- **Automated suggestions for model construction and refinement.** We are designing and implementing a mechanism to learn how factors interact in non-trivial ways in scientific models. In particular, we are developing a method for learning disjuncts, conjuncts, and negations. This mechanism starts with the model developed by the scientist user and make recommendations for possible adjustments to make it more complete by performing statistical data mining and machine learning.
- **Data validation and repair.** Recognizing that data contains errors is plausible once we understand the relationships between data sets. That is, if we are able to develop models of the correlations between sets of data, then we can build systems that notice when these correlations do not hold in new data, indicating possible errors in data. For instance, if we know that public sentiment tends to vary similarly between nearby towns, then when one town shows anomalous behavior, we can reasonably suspect problems with the data. There might be local issues that cause the anomaly, but it is, at least, worth noting and bringing to the attention of the user of the data and model. As MAT is designed to help analyze models and recognize inter-data relationships, it is primed to perform exactly this analysis. Existing methods perform similar types of analysis for environmental data (Dereszynski & Dietterich, 2007, 2011). For instance, a broken thermometer can be identified and the data from it even estimated by looking at the temperature readings of nearby thermometers, which will generally be highly correlated.
- **Application to multiple scientific modeling domains.** To ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains, we are identifying and building at least one neurological and/or physiological model and analyzing the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model. The initial MAT effort focused on HSCB models; by focusing this effort on harder-science models at much shorter time durations, we believe we can effectively evaluate an interesting range of applications of the MAT tool.

3. Current Activities and Status

During the current reporting period, we made progress on a number of fronts, including improved support for describing temporal constraints on causal models, improved support for visualizing the validation results of the causal model analysis, and a number of usability improvements to support our users. We also wrote and submitted a paper and paper abstract on MAT and continued to support a variety of users and programs that are using MAT.

We describe our progress during the current reporting period for each of these areas in turn.

3.1. Causal Models

In this section, we begin my summarizing a number of methods for validating causal relationships in models. Such validation can help researchers produce more robust models of

complex systems, facilitating testing of dependencies that might otherwise be missed, assumed away, or taken for granted. These approaches have all been described to various levels of detail in previous reports; in this report we summarize the current state.

3.1.1. Granger causality for validating dependencies

Granger causality was originally introduced for economic models (Granger, 1980, 1969) to help deal with the problem of *temporal offsets* described above. It can, however, be adapted as a validation test for causality in socio-cultural data. Granger causality makes two assumptions: (1) the effect does not precede the cause, and (2) the causal variable provides information about the effect that would otherwise be unavailable.

Definition 1 (Granger Cause). *The temporal variable X* Granger causes *temporal variable Y iff* $P(Y_t \mid Y_{t-1}^{t-L}) \neq P(Y_t \mid Y_{t-1}^{t-L}, X_{t-1}^{t-L})$ *where L is the maximum time lag,* a_i, b_j *are parameters in a linear combination,* ϵ_1, ϵ_2 *are error terms, and*

$$P(Y_t \mid Y_{t-1}^{t-L}) = \sum_{l=1}^{L} a_l Y_{t-l} + \epsilon_1 \quad (1)$$

$$P(Y_t \mid Y_{t-1}^{t-L}, X_{t-1}^{t-L}) = \sum_{l=1}^{L} a_l Y_{t-l} + \sum_{l=1}^{L} b_l X_{t-l} + \epsilon_2 \quad (2)$$

A variable *X* is a Granger cause of *Y* if *Y* can be better predicted using the histories of *X* and *Y* than just of *Y* alone. We can validate this relationship through hypothesis testing. If equation (2) is statistically more accurate than equation (1) using an *F* statistic, then a causal relationship between *X* and *Y* is valid. This test is given in Figure 1.

Granger Causality *F*-test Validation Procedure
H_0: *X* does not Granger cause *Y*.

1. Choose the desired significance level α and identify the critical value *c*.
2. Use equations (1) and (2) to compute the *F*-statistic where (1) is a restricted model of (2) such that all of the coefficients $b_{t-1} = b_{t-2} = \ldots = b_{t-L} = 0$.
3. If $F > c$, then the null hypothesis can be rejected and a causal relationship is validated.

Figure 1. Granger causality *F*-test validation

3.1.2. Dynamic time warping for uneven temporal relationships

Many causal relationships are imperfectly represented in the observed data. This is particularly salient in complex socio-cultural systems where variability in human behavior produces uneven temporal delays between cause and effect. For example, lower employment rates may cause an increase in crime anywhere from 6 to 12 months in the future. Such relationships cannot be captured by standard statistical analyses, including Granger causality, which assume a stationary process with a consistent time lag. To validate these relationships, we borrowed and extended the dynamic time warping (DTW) algorithm from gait recognition (Salvador & Chan, 2007; Myers & Rabiner, 1981), where DTW is used to identify a gait from two motion curves even when a person speeds up or slows down. The DTW algorithm compares the two time series to find the optimal alignment by “warping” one series, i.e., stretching or shrinking it along its time axis.

Definition 2 (Warp Path). *Given two time series X and Y of size n and m a* warp path *W is a sequence W=w_1, w_2, …, w_K where K is the length of the path and each element w_k = (i, j) represents a mapping between point i in X with point j in Y. The optimal warp path minimizes the sum of distances between the mapped points*

$$argmin\ Dist(W) = \sum_{k=1}^{k=K} Dist(w_{ki}, w_{kj})$$

where $Dist(W)$ is the distance of warp path W and $Dist(w_{ki}, w_{kj})$ is the distance between point i in series X and point j in series Y.

A warp path identifies a series of cells in an $n \times m$ two-dimensional cost-matrix D, where each $D(i,j)$ is the minimum distance warp path that can be constructed from X up to x_i and Y up to y_i. The final entry contains the optimal path over the full series.

Because causality can only impact the future, we have expanded DTW to handle the one-directional case in a new algorithm ForwardDTW. Rather than matching points in both directions, ForwardDTW only matches the points in X with future values of Y, i.e., the entry $D(i,j)$ is only computed for $i < j$. ForwardDTW now allows us to use DTW to validate causal relationships—the smaller the warp distance between X and Y, the stronger the causal link. A user can specify this causality threshold to determine when a relationship will be considered validated as shown in Figure 2.

Some advantages of DTW over other time series analyses are that it can account for missing data and compare series with different time scales or sampling frequencies. DTW is also very visual, making the results easy to interpret by human analysts.

Dynamic Time Warping Validation Procedure

1. Set a warping threshold *t*.
2. Use the ForwardDTW algorithm to compute $min\ Dist(W)$ for X and Y.
3. If $min\ Dist(W) < t$, a causal relationship is validated.

Figure 2. Dynamic time warping validation procedure.

1.1 Convergent cross mapping for dynamic feedback models

Granger causality and DTW can identify causal relationships and consider complex temporal factors. However, many social systems contain feedback relationships, where dependency between variables is bi-directional—declining economic output may increase levels of political violence, which further depress the economy, etc. Such relationships are extremely difficult to validate using standard approaches. To analyze cyclic causality, we utilize the convergent cross mapping (CCM) algorithm(Sugihara, May, Ye et al., 2012). CCM was first introduced in biology to model predator/prey systems, but can be adapted to model the interrelationships in other types of scientific data.

To use CCM, we derive a set of vectors for variables X and Y called the "shadow manifolds" that represent a topological projection of the underlying dynamic system.

Definition 3 (Shadow Attractor Manifold). *For a time series variable X the* shadow attractor manifold M_X *consists of points* $x(t) = (X(t), (X(t-\tau), X(t-2t), \ldots, X(t-E\tau))$ *where τ is a sampling time lag and E is the manifold dimension.*

For subsets of the time series X and Y of length L we can construct manifolds M_X and M_Y. CCM will then determine how well local "neighborhoods"—small regions of M_X —correspond to neighborhoods in M_Y. If X and Y are causally linked, there will be a one-to-one mapping between points in M_X and M_Y. To compute this cross mapping, we use a neighborhood in M_X to predict the values of contemporaneous points in M_Y and compute the correlation ρ between the predicted values $\hat{Y}(t)$ and the real values $Y(t)$. If a causal relationship exists, predictions of the state of Y from X (and vice versa) will improve asymptotically as the amount of data (L) increases, i.e., the mapping of X and Y will converge to perfect predictability $\rho = 1$. Fig. 3 shows this process.

Convergent Cross Mapping Validation Procedure

1. Randomly choose segments of length L from X and Y.
2. Construct shadow attractor manifolds M_X and M_Y for X_L and Y_L.
3. Compute the cross mapping ρ between M_X and M_Y in both directions.
4. If ρ converges toward 1as L increases, then there is a causal link.

Figure 3. Convergent cross mapping validation procedure

3.1.3. Qualitative pattern analysis

To aid social scientists in validating models that are more readily described as patterns of qualitative features than as quantitative relationships, we use a two-part mechanism that (1) identifies features, and (2) searches for patterns among those features. For instance, we would first search through an economic data set to extract periods of "economic instability" and a crime database to extract periods of "increasing crime." These feature definitions are themselves quantitative and can be described directly using mathematical formulae acquired through a process of "learning-by-demonstration", where our system infers mathematical descriptions of subsets of the data that are labeled by the scientist, or these features can be algorithmically extracted by automatically looking for "interesting" subsets of data.

To support the latter two mechanisms, we have developed a structural featurization "language" using six common morphologies shown in Figure 4 from (Olszewski, 2001). A time series can be represented as a sequence of these components.

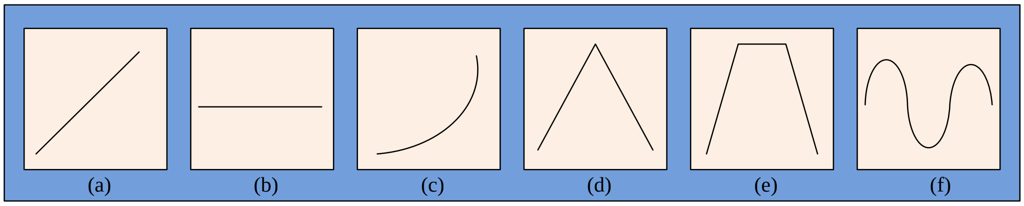

Figure 4. Features (a)slope, (b)constant, (c)exponential, (d)triangle, (e)trapezoidal, (f)sinusoidal

In the case of learning-by-example, our algorithm finds the best match between the user-identified data to these feature types (and associated parameters). In the case of automatically extracting "interesting" features, we identify features with high explanatory value that are characteristic of a time series, but not so common that they are meaningless. We developed a variant of the tf-idf algorithm (Manning, Raghavan, Schutze et al., 2008), typically used in search engines, to "score" features. The (-*log tf*) in the definition penalizes features that occur too often.

Definition 4 (Feature Score). *Given a database D and feature f from time series Y in D, the* feature score *is the function:*

$$tfidf\_score = tf \times \frac{-\log tf}{df}$$

where tf is the density of feature f in series Y and df is the density of the feature in D.

Once features are identified, we can search for causal patterns of features that are posited by the model. For instance, if the model claims that increases in crime (*C*) follow periods of economic instability (*E*) (within some time window), we can search for cases where instances of *C* are not preceded by *E*s, where *E*s are not followed by *C*s, and where the specified pattern is actually found in the data.

In addition to these techniques for analyzing causality, we also made progress on software tools and visualization implementations to support the definition and validation of causal models that leverage the techniques just described. In the next sections, we describe these software improvements.

3.1.4. Graphical Tool for Specifying Constraints on Causal Relationships

One of the challenging issues in describing causal models is how to make it possible (and, ideally, easy) to specify the temporal constraints on a causal link. For instance, if the hypothesis is that A (say, unemployment) leads to B (say, crime), and the data indicates that there is a jump in crime after a spike in unemployment, the temporal distance between A and B is still a significant factor in whether we want to associate the two events. If B happens 6 months after A, it could be counted as evidence. It B happens 6 decades after A, we probably want to treat them as unrelated events. Specifying these temporal constraints is even more difficult, though, in that the relationships are not simply "starts-after," but can relate to other factors, such as

when A or B ended. We have previously described a mechanism for specifying these constraints, but users found in unintuitive and wanted additional functionality added, so we devoted time during the current reporting period to updating and expanding this functionality.

To make the creation of causal models as easy as possible, a new visualization of the causal relationships has been developed. Now, the user can see how the cause and effect are compared when validating a causal relationship. For example, in the causal model in the following screenshot, Concept A causes Concept B using the simplest of causal relationships where Concept A must start sometime before Concept B.

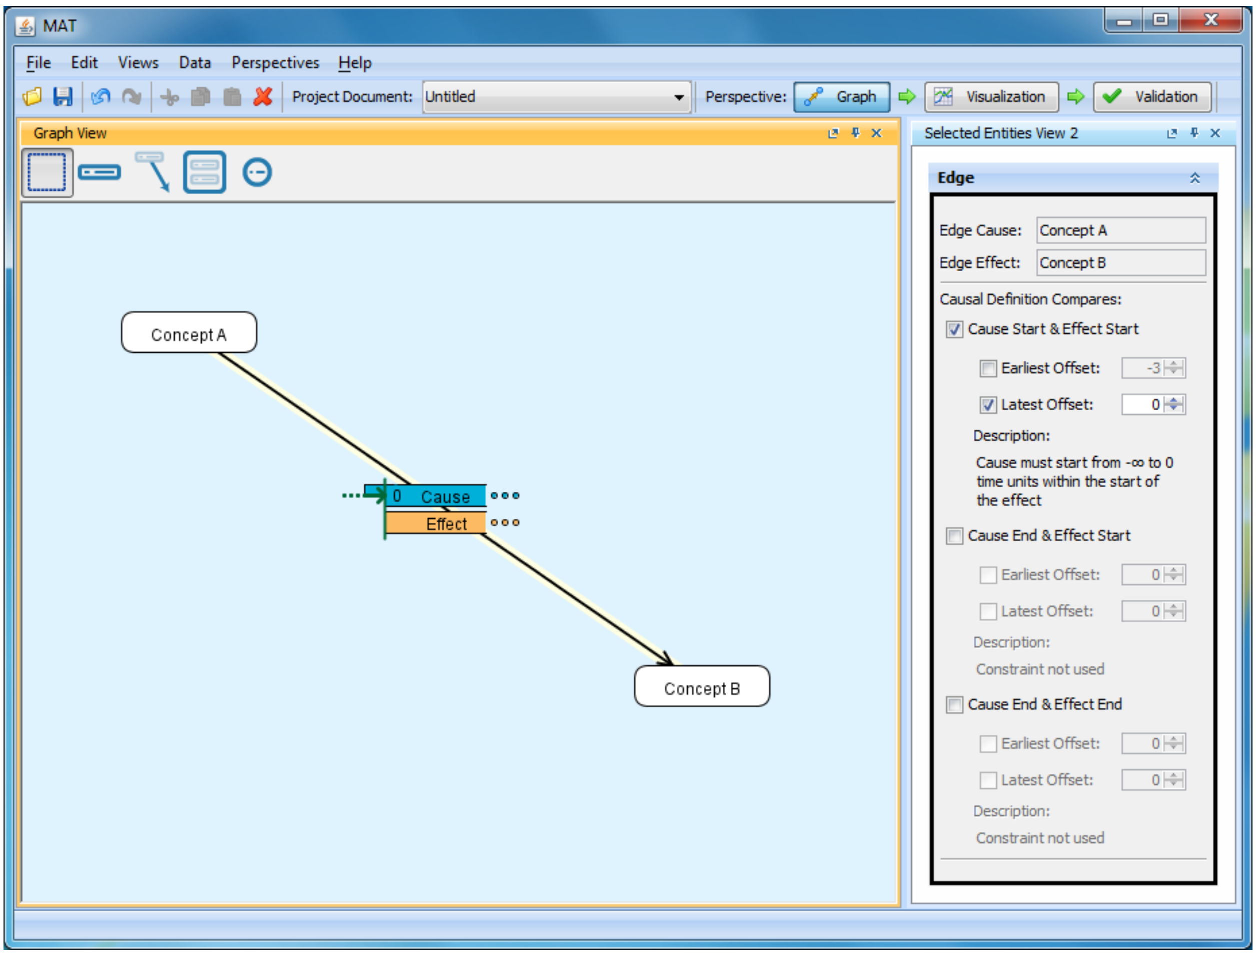

However, this is a lenient causal relationship since Concept A could have started and ended many years (or centuries) before Concept B. Therefore, a more useful causal relationship would include a window that restricts when the cause starts relative to the start of the effect. Also, a second type of constraint can be included that restricts when the cause ends relative to the end of the effect. The following screenshot shows this new causal relationship.

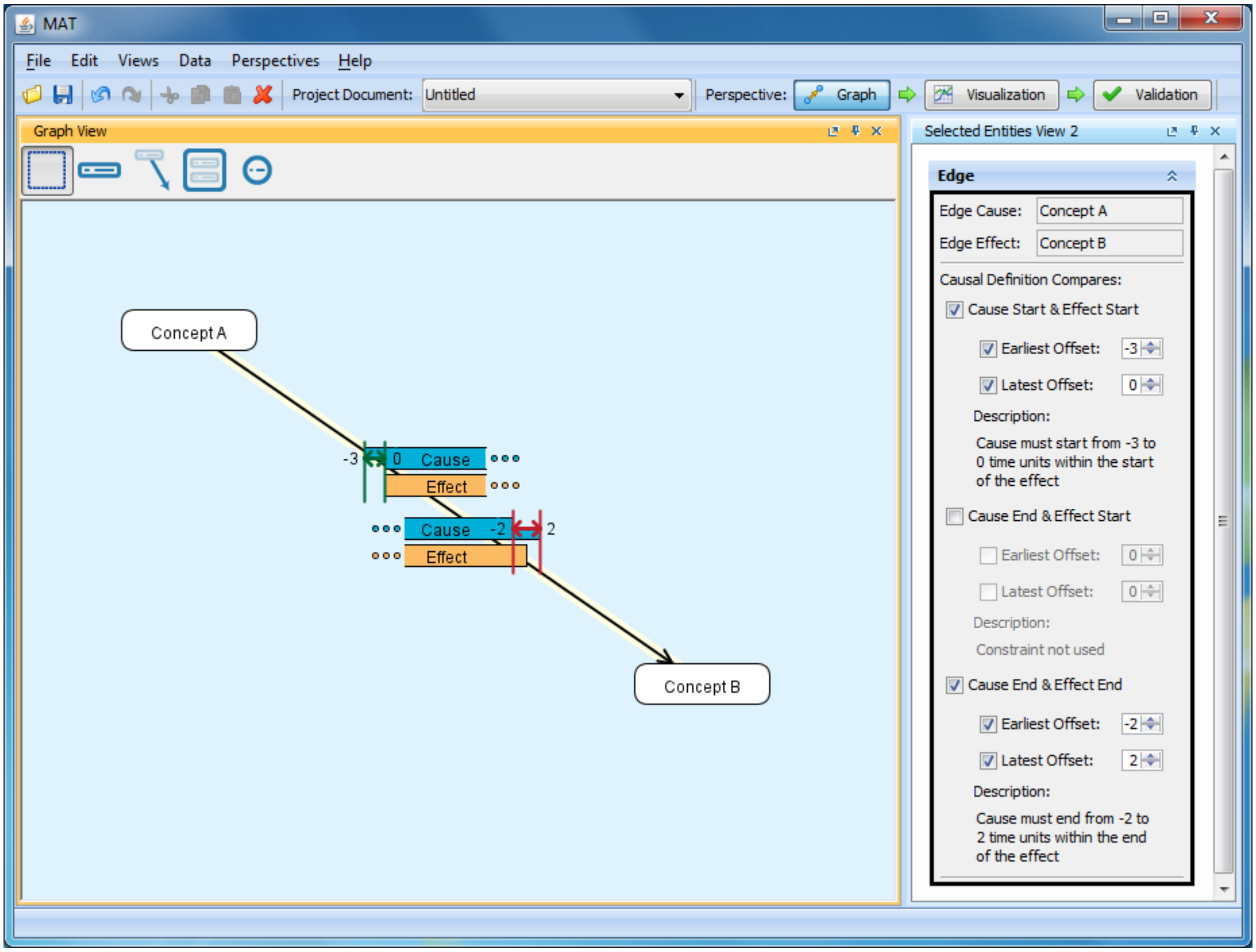

The panel on the right allows the user to add, remove, and modify three types of causal constraints that compare: (1) cause and effect start, (2) cause and effect end, and (3) cause end and effect start. Also, the panel on the right provides a verbal description of each of the causal constraints to further improve the user's understanding of the causal relationship.

3.1.5. Graphical Tool for Visualizing Causal Relationship Validation Results

The validation visualization shows whether the causal model was supported by the empirical data. In the following screenshot, the green events are effects that are supported by a cause.

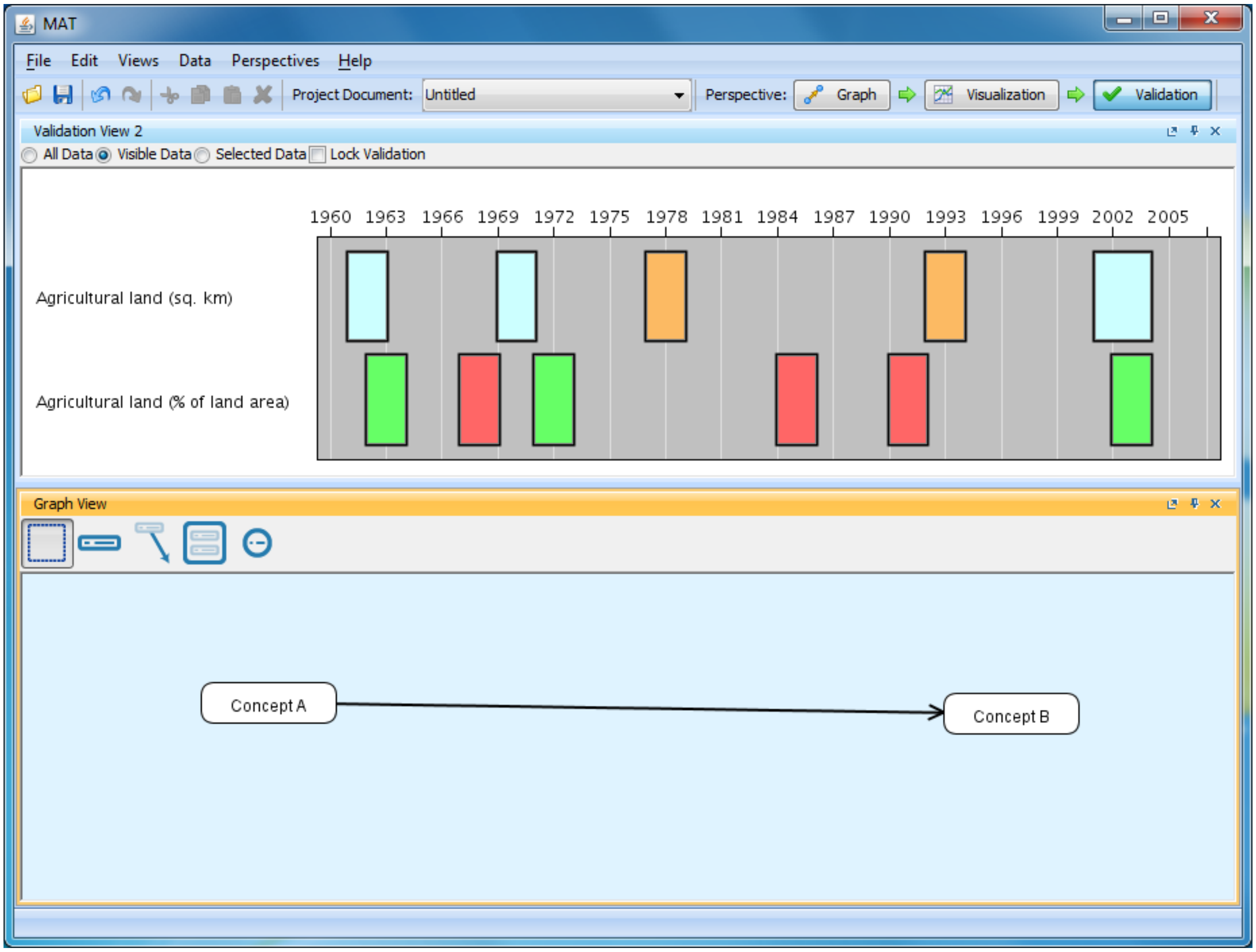

In simple examples, it is easy to tell which cause supports each effect, but in more complex causal models this may not be so clear. Therefore, the visualization can now include arrows showing the valid causal relationships. For example, in the following screenshot, the user selected the causal relationship in the causal model, which causes the valid examples of that causal relationship to be displayed as arrows pointing from cause to effect. This allows the user to easily and quickly see why an effect was supported.

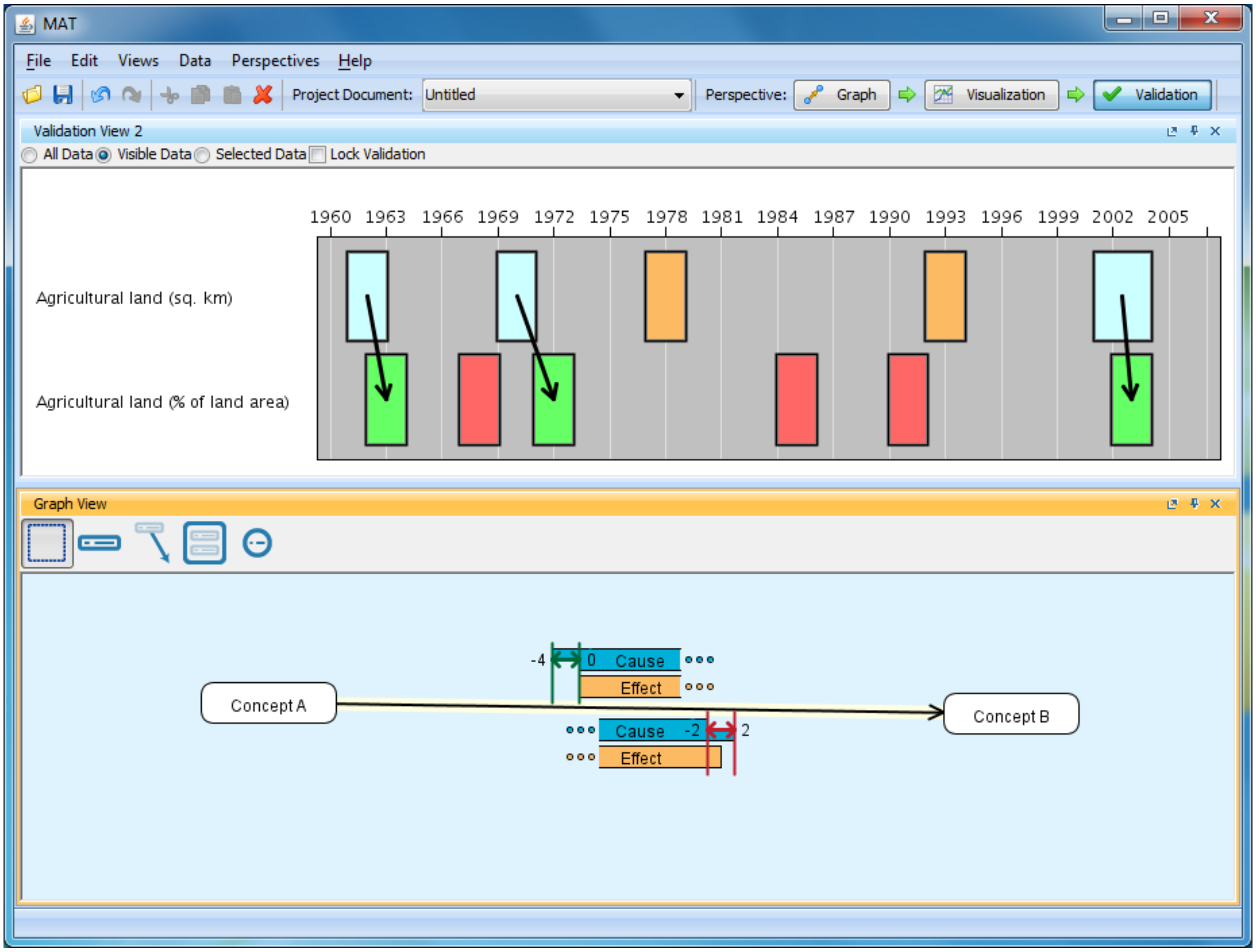

3.2. Continued Debugging, Feature Enhancement, and Release Preparation

In preparation for the Version 7 release of MAT we continue to correct a range of minor bugs and small functionality lapses. The underlying infrastructure of the application is being improved to handle a more expansive set of use cases and user environments. In particular, when MAT was first built, the intended audience was social scientists working with recent socio-economic time series data. The newest version of MAT will support a much wider range of potential users including physical scientists working in non-specific, relative, or extreme time scales.

The new version will allow them to not only work with existing data series in fixed formats, but will allow them to easily import a wide range of time series data formats and, more importantly, to modify and manipulate data as part of their model. We found this capability to be very important because in many cases model data comes from different sources and is not scaled in a mutually compatible way. Also, different data series might be in different units. Before, the scientist would have to manually massage their data to make it compatible before using it in MAT. Now, they can integrate such manipulations directly into their project.

In summary we worked on the following items for the new version:

- Greatly improved flexibility by using a new universal time-data type
- Improvements to the data importer to support new universal time-data type
- Rework of MAT internals to support use and validation of universal time-data type
- Usability improvements for analytical popups (e.g., Grainger causality, correlations)
- Addition of the new data synthesis capability
- Automated feature discovery
- Added a feature editor that enables the user to edit features in a table instead of having to use the graphical editor previously provided. This can be much more efficient in some cases (e.g., where there are multiple features to edit, where the quantitative details of the features are known). This new features is shown in Figure 1.

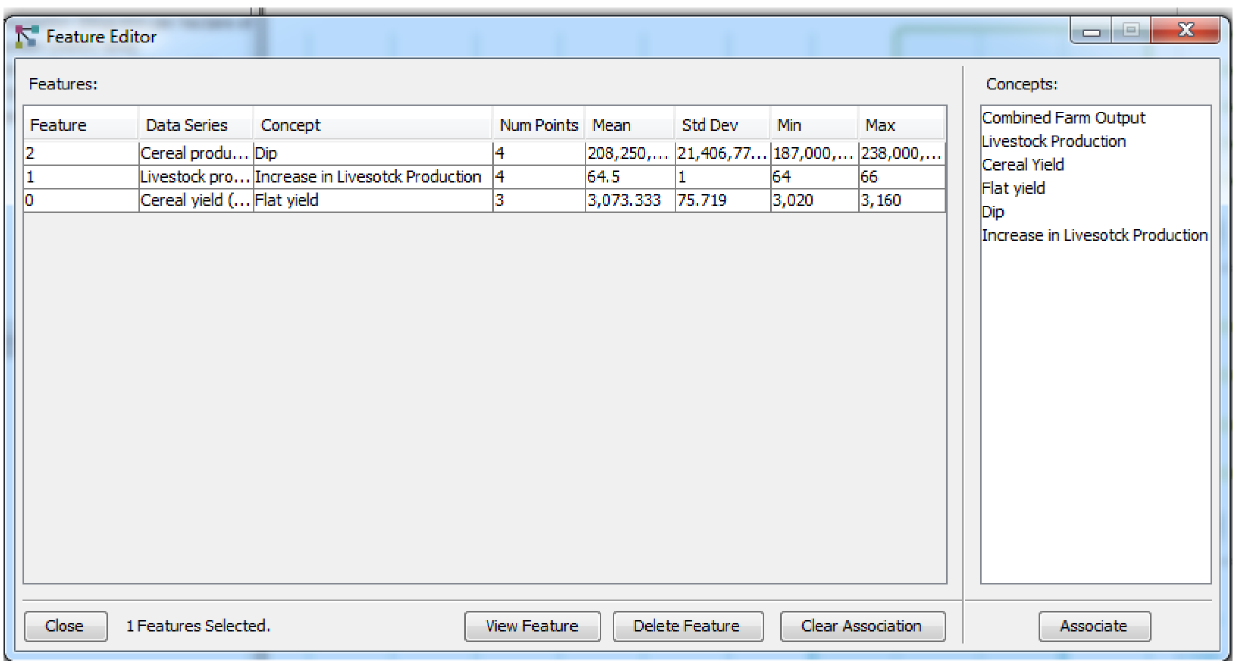

Figure 5. New Feature Editor in MAT

- Improved the usability of the model editing interface
- Other bug fixes and corrections

Along with the usability and productivity improvements we are also mapping out the next generation of functionality in MAT. The ability to perform automatic relationship discovery is one advanced feature which we are working on including in the new version at some level. This feature allows the user to select multiple data sets and suggest possible relationships between

them. The system will then automatically generate model nodes to express the discovered relationships.

4. Planned Activities

During the upcoming reporting period, we plan to release a new version of MAT that includes support for:

- Data synthesis
  - new tab with draw interface (analogous to modeling tab)
  - ability to add synthesized data to loaded data sets
  - ability to select and use loaded data sets as input
- Data type upgrade
  - all times represented by new object structure
  - supports any kind of time or date including intervals
  - upgrade to data import to use new time data type
  - upgrade to visualization to use new time data type
  - upgrade to popup analysis dialogs to use new time data type
  - compatibility adjustment to time constraints/validation to use new time data type
- Feature Discovery, stabilization and functionality signoff
- Causality popup stabilization and functionality signoff

We also hope to complete work on the automatic feature detection algorithm and implementation and to continue work on the next release of MAT, which will include support for:

- Improved causality analysis and reporting
- The new time-constraint interface
- The new validation visualization interface
- A new, much-more-flexible windowing layout system and cross-pane data sharing (this will, for instance, support highlighting of features in the data-visualization pane when they are chosen in the validation tab)

5. Evaluation and Transition

We continue to focus on making MAT available to the government and academic research communities and to look for opportunities to use MAT on a variety of ongoing research efforts. Table 1 summarizes our progress in this regard to date. We will continue to update this table as we make additional progress and will include it as a regular part of future status reports.

Program	**Customer**	**Comments**
On-going efforts		
Tourniquet Master Trainer (TMT) (Phase I SBIR)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	MAT is being used to visualize and analyze data from sensors on a medical manikin that indicate whether a number of novel medical devices used to combat junctional and inguinal hemorrhaging are being applied properly. This program is about to begin a Phase II where MAT will continue to be used both by Charles River Analytics and our partners at the University of Wisconsin.
Laparoscopic Surgery Training System (LASTS) (Phase II SBIR)	US Navy's Office of Naval Research (ONR)	Under lasts, Charles River and Caroline Cao at Wright State University are using MAT to analyze data collected from the location of the laproscopic surgery tools tools during an experiment. Surgical tools are instrumented with markers and 3D data is collected on their location as the person performs the task. This is an ongoing Phase II SBIR program.
Cognitive Readiness Agents for Neural Imaging and Understanding Models (CRANIUM) (Phase I SBIR)	US Navy's Office of Naval Research (ONR)	MAT was used to visualize and extract patterns of stress and workload from neuro-physiological data for training systems. This was a Phase I SBIR program that did not progress to Phase II.

Business Intelligence Visualization for Organizational Understanding, Analysis, and Collaboration (BIVOUAC) Phase II SBIR	US Navy's Space and Naval Warfare Systems Command (SPAWAR)	MAT is being evaluated as part of the BIVOUAC SBIR program, which provides data analysis and visualization for Enterprise Resource Planning (ERP) systems for the Navy. This is an ongoing Phase II SBIR program.
Adaptive toolkit for the Assessment and augmentation of Performance by Teams in Real time (ADAPTER) (Phase I SBIR)	US Air Force Research Lab Human Effectiveness Directorate (AFRL/RH)	MAT is being used to analyze neuro-physiological data from cyber operators to evaluate cognitive workload during team-based cyber operations. This is an ongoing Phase I SBIR program. A Phase II proposal has been requested and is currently being written.
Anticipated Efforts		
Enhancing Intuitive Decision Making Through Implicit Learning (I2BRC) (ONR Basic Research Challenge BAA)	US Navy's Office of Naval Research (ONR) Charles River is a subcontractor to DSCI MESH Solutions, LLC	The intention is to use MAT to help analyze neuro-physiological data to help better understand how implicit learning and intuitive decision making work. This is an ongoing BAA program, though no data has yet been collected to analyze.
A system for augmenting training by Monitoring, Extracting, and Decoding Indicators of Cognitive Load (MEDIC)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	We are evaluating the practicability of using MAT to analyze and visualize neuro-physiological data from combat medic trainees to identify periods of stress and cognitive overload. This is a recently started SBIR Phase I program in its initial design and exploration phase.

Table 1. MAT Transition and Use Progress

In addition we have provided copies of MAT to the following institutions based on their requests for the software: the University of Michigan, Arizona State University, Kansas State University, University of California at Los Angeles, the Naval Medical Research Unit at Wright Patterson Air Force Base, Concordia University (Montreal), the University of Wisconsin, and the Air Force Research Laboratory's Human Effectiveness Directorate.

Finally, we submitted a paper on the various causality analysis techniques in MAT to the Social Computing, Behavioral Modeling and Prediction 2014 Conference and a paper abstract on using MAT for data-driven model refinement and validation to the annual meeting of the American Political Science Association.

6. Budget and Project Tracking

As of October 31, 2013, we have spent $439,813, or 47% of our total budget of $928,224, in 47% of the scheduled time. Our current funding is $444,945.00, so we have spent 99% of our available funding. [Note: an additional funding increment was received during the writing of this report, bringing our total funding to $524,945.]

We believe we are in good shape to complete the project on time and on budget.

7. References

Dereszynski, E. W.& Dietterich, T. G. (2011). Spatiotemporal models for data-anomaly detection in dynamic environmental monitoring campaigns. *ACM Transactions on Sensor Networks (TOSN)*, 8.

Dereszynski, E.& Dietterich, T. (2007). Probabilistic models for anomaly detection in remote sensor data streams.

Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424-438.

Granger, C. W. (1980). Testing for causality: a personal viewpoint. *Journal of Economic Dynamics and Control*, 2, 329-352.

Manning, C., Raghavan, R., and Schutze, H. (2008). *Introduction to Information Retrieval*.: Cambridge University Press.

Myers, C.& Rabiner, L. (1981). A level building dynamic time warping algorithm for connected word recognition. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 29, 284-297.

Neal Reilly, S. (2010). *Validation Coverage Toolkit for HSCB Models*. Charles River Analytics Inc.

Neal Reilly, W. S., Pfeffer, A., and Barnett, J. (2010). A Metamodel Description Language for HSCB Modeling. In *Advances in Cross-Cultural Decision Making*.

Neal Reilly, W. S., Pfeffer, A., Barnett, J., Chamberlain, J., and Casstevens, R. (2011). A Computational Toolset for Socio-Cultural Data Exploration, Model Refinement, and Model Validation. In *Human Social Culture Behavior (HSCB) Focus*.

Olszewski, R. (2001). *Generalized feature extraction for structural pattern recognition in time-series data*. Carnegie Mellon University, School of Computer Science.

Salvador, S.& Chan, P. (2007). Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11, 561-580.

Sugihara, G., May, R., Ye, H., Hsieh, C., Deyle, E., Fogarty, M., and Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338, 496-500.